Mutual Initiative in Human-Machine Teams

David J. Bruemmer, Julie L. Marble, Ph.D., and Donald D. Dudenhoeffer

***Abstract*-- Autonomous systems are widely used today in industry. The human-machine relationship in these systems is primarily that of a human supervisory role. This paper explores the concept of human-robot teams where each member of the team has the ability to assume initiative within a task. Key to this effort is not only the ability of the human to understand and predict robot performance, but the robot's ability to identify human needs and select intervention points to assume different levels of initiative. The objective is to incorporate mobile autonomous robots into human teams to augment both the human's cognitive and physical abilities in the performance of potentially hazardous tasks.**

***Index Terms*—Automation, Cognitive science, Human factors, Intelligent robots, Mobile robot dynamics, Robots.**

I. INTRODUCTION

This paper discusses the concept of human-robot teaming in joint task accomplishment, specifically with regards to the ability of both the human and robot to take initiative in task performance. The role of a robot in such a situation is to augment human cognition and physical activity. The objective of such teaming is to increase the level of task performance by leveraging off the unique capabilities of each performer. Such performance improvements may affect time to complete work, the ability to conduct the task, and efficiency. The robots for this teaming are not remotely controlled manipulators, but posses a level of autonomy, i.e., they are able to act without direct human interaction at some level.

The use of a robot in this particular scenario can be likened to that of a police dog and the human partner. The dog provides the human augmented capabilities such as sniffing to find drugs or tracking a fugitive through the woods. Additionally, the dog augments the human's physical capabilities by searching human-inaccessible areas and the ability to confront a would be threat at a standoff distance from its human partner.

Within the Idaho National Engineering and Environmental Laboratory (INEEL), as well as other Department of Energy (DOE) complexes, exist extremely hazardous environments. These environments may contain high radiation areas, radioactive contamination, hazardous material contamination, or a combination of all the above. Worse yet, the levels of such contamination and the associated human hazards may not be fully known prior to entering the area.

Human-robot teams offer a means to accomplish tasks within such areas efficiently while minimizing the hazard to the human element. In this example, the robot can augment human cognition in sampling and analyzing changing environmental conditions to assess human risk and stay time. Additionally, the robot can carry equipment, remotely position sensors, and conduct physical labor as a means of augmenting physical activity.

The use of autonomous systems is not a new area of research. To a large degree, however, such human-machine systems have not explored the teaming concept, but have utilized a human supervisory control schema. Specifically, supervisory control of a process implies that a human operator communicates with the machine to gain information and issue commands while the machine implements these commands through artificial sensors and actuators to control the task or process [1]. Within this control schema, the machine may exhibit different levels of autonomy in the task.

Sheridan and Verplank [2], [3] introduced a scale to describe levels of Human-Machine Interaction for the accomplishment of the process within the supervisory control system. These levels of interaction are:

1. Whole task done by human except for actual operation by machine;
2. Human asks computer to suggest options and selects from the options;
3. Computer suggests options to human;
4. Computer suggests options and proposes one of them;
5. Computer chooses an action and performs it if human approves;
6. Computer chooses an action and performs it unless human disapproves;
7. Computer chooses an action, performs it, and informs human;
8. Computer does everything autonomously.

It is important to note that the levels of autonomy and thus the levels of interaction relate to the machine's cognizance of the process and the environment surrounding the process. This does not address the aspect of the machine's recognition of the human's performance within the process. This is a subtle, yet very important distinction.

In the case of human-robot teaming, the paradigm must shift. Not only must the robot focus on the task, but additionally, the robot must identify the needs and shortcoming of the human element.

This work is supported through the INEEL Long-Term Research Initiative Program under DOE Idaho Operations Office Contract DE-AC07-99ID13727.

D. Bruemmer, J. Marble, and D. Dudenhoeffer are with the Idaho National Engineering and Environmental Laboratory, PO Box 1625, MS 3605, Idaho Falls, ID 83415. (email: bruedj, marbjl, dudedd@inel.gov)

II. TEAMING

The Sheridan and Verplank levels of Human-Machine Interaction still serve as a basic model for the incorporation of autonomous and semi-autonomous systems into mixed teams. However, this taxonomy of interaction relegates the machine to the role of a subordinate and assumes that the machine cannot override human commands or take initiative to assign tasks to the human.

We submit that rather than conceive of machines as mere tools, wholly subordinate to humans commands, it is more effective to consider the machine as part of a dynamic human-machine team in which each member is invested with agency – the ability to actively and authoritatively take initiative to accomplish task objectives. Within this schema, each member has equal responsibility for performance of the task, but responsibility and authority for particular task elements shifts to the most appropriate member, be it human or machine.

There are several possible roles for each member of a mixed human-machine team: Tool (applied to a single task with direct control by a supervisor); Subordinate (member is given a high level task with less direct interaction by a supervisor); Equal (no direct supervision; each team member is wholly responsible for some aspect of the task); or Leader (directs the other members to perform a high level task). Note that these levels of interaction imply that the machine aspect of the system could direct the human components of the system to perform aspects of the task; that is, it is reasonable that the machine may lead at least some aspect of the task where its supervision would be more effective than the reverse.

Categories for characterizing human and robot initiative can provide a skeleton around which to design robotic control architectures. In addition, they provide a basis for user understanding and trust. The greatest challenge to the adoption of autonomous robots is that of trust.

However, while it is profitable to categorize levels of initiative, it is important to understand that robots for targeted applications will most likely have modes of control that do not fit cleanly into one category. Optimal task performance may require that the robot assume a "leader" position over one aspect of the task and yet in all other respects place the robot in a subservient position.

The robot is often in a much better position than the human to react to the local environment, and consequently, the robot may take the leadership role regarding navigation. As leader, the robot can then "veto" dangerous human commands to avoid positive and negative obstacles. Robot intelligent perception however is still very immature and when faced with half-open doors or empty boxes, the superior knowledge of the human requires a role change and leadership returns to the human for such judgment decisions. These changing roles implicate the importance of considering human factors in the design of human-robot teams and mutual initiative.

Given the desire to employ robots in hazardous, critical environments, the ability to shift a robot in and out of the leadership role presents a conundrum. The user comes to rely on the self-protective capabilities of the robot and yet, at times, must override them to accomplish a critical mission.

For instance, when faced with an unknown box obstructing the path, the user may shift the robot out of the leadership responsibility for navigation, but grant the robot the "right" to refuse human commands when the physical resistance to motion is beyond a certain threshold. For other tasks, the user may need to drive the robot to where it is touching an obstacle. The user can curtail the robot's collision avoidance initiative and yet customize a "last resort" channel of initiative based on bump sensors and short-range infrared break beams.

Ideally, we need control systems that allow the user to configure particular "channels of initiative" that crosscut broad categories. The roles of each team member are bounded by a complex and changing web of capabilities and limitations to which each member must adapt and respond. The ability of the human to develop accurate understanding of robot behavior is essential if this adaptive role switching is to work effectively. One of the most fascinating areas of future work is the need for the robot to be imbued with an ability to understand and predict human behavior. This is the focus of on-going work at the INEEL, which is discussed in a later section.

III. TEAM NEED'S AND EXPECTATIONS

The fundamental aspect of a human team that distinguishes it from a simple group is the presence of a shared goal. Furthermore, effective teams typically anticipate the information needs of teammates via a shared mental model of the task and current situation [4]. In many teams, members do not fulfill a single, static role. In fact, many members end up playing roles distinct from the task skills that gained them admission into the group. Within a healthy team, group roles evolve dynamically to meet new and unforeseen challenges.

When we apply team role theory to human-robot interaction, we find that very few true human – robot teams exist. Machines are generally used as tools, rarely as peers and scarcely ever as task leaders. Just as roles evolve between human team members, mixed human – robot teams should be able to modulate the level of robotic initiative in order to balance changes in the environment, task, and capabilities of other team members.

The ability to transition roles and responsibility among the members of mixed human-robot team presents new challenges. After all, the dynamics of performance for robots are drastically different than for humans (i.e., communication, perception, degradation, etc.) Performance of a robot may far exceed human abilities in certain elements of a task; however, robots are notoriously unable to degrade gracefully in the face of component failure or unforeseen changes in the environment or task. Thus, while robots are often more reliable than humans within certain parameters, they ultimately lack the reliability necessary for tasks where such parameters cannot be fully guaranteed. The answer to this challenge is not to keep robots as passive tools. Rather, we must have a team dynamic where roles can shift to compensate for the unique kinds of failures and limitations possessed by each member.

This need for dynamic role switching makes critical the consideration of human factors in the design of the human-machine interface, especially those issues regarding how each member should communicate his, her, or its role in the team. This includes the need for each team member to be able to determine and be aware of the current capabilities of other team members. Among human teams, members are often able to identify when another member is becoming overwhelmed or over worked and task performance is degraded. It is more difficult for a human to realize when a machine is overworked and more difficult still for a machine to recognize when a human is overwhelmed or overworked.

A. Theory of Robot Behavior

The need for human and robot to predict and understand one another's actions presents a daunting challenge. Indeed, the challenge of supporting the situation awareness of the user has occupied the field of robotics for decades. For each level of robot initiative, the user must develop a unique set of expectations regarding how the robot behaves, that is, an understanding or theory of the system's behavior, here after referred to as a theory of robot behavior (TORB). By TORB we mean that the human operator is able to quickly and accurately predict:

1. Actions the robot will take in response to stimuli from the environment and other team members;
2. The outcome of the cumulative set of actions.

The human may acquire this theory of behavior through simulated or real world training with the robot. Most likely, this theory of behavior will be unstable at first, but become more entrenched with time. Further work with human participants is necessary to better understand the TORB development process and its effect on the task performance and user perception. Recalling the example of a police working dog, the policeman and his canine companion must go through extensive training to build a level of expectation and trust on both sides. Police dog training begins when the dog is between 12 and 18 months old. This training initially takes more than four months, but critically, reinforcement training is continuous throughout the dog's life [5]. This training is not for just the dogs benefit, but serves to educate the dog handlers to recognize and interpret the dog's movements which increases the handler's success rate in conducting task. It must also be noted that no two dogs will act exactly the same.

However, in our research we are not concerned with developing a formal model of robot cognition, but rather require that the human understand and predict the emergent actions of the robot, with or without an accurate notion of how intelligent processing gives rise to the resulting behavior. When faced with a robot that can orchestrate task elements, the critical issue will not be how the robot or machine "reasons," but rather whether the human team members can accurately predict robotic responses and understand how cumulative actions and responses converge to fulfill task objectives.

Many applications require the human to quickly develop an adequate TORB. One way to make this possible is to leverage the knowledge humans already possess about human behavior and other animate objects, such as pets or even video games, within our daily sphere of influence. For example, projects with humanoids and robot dogs have explored the ways in which modeling emotion in various ways can help (or hinder) the ability of a human to effectively formulate a TORB [6].

Regardless of how it is formed, an effective TORB allows humans to recognize and complement the initiative taken by robots as they operate under different levels of autonomy. It is this ability to predict and exploit the robot's initiative that will build operator proficiency and trust. The development of a theory of robot behavior will also allow the user to switch between and configure the robot's levels of initiative to suit the needs and components of the task at hand. Informal experimentation with our own adjustable autonomy control system indicates that no one role is, in and of itself, optimal for the robot or human. Rather, the human – robot team is best enabled to accomplish a complex task when given the opportunity to shift back and forth between modes of control. Using the real world robot system, the INEEL is in the planning stages of performing human subject experiments to assess how overall performance is affected by allowing the user to shift roles for the robot on the fly.

B. Theory of Human Behavior

In the 1940's and 1950's, Isaac Asimov through a series of short stories developed the three Laws of Robotics that governed robotic behavior throughout his books [7]. In 1985, Asimov added an additional law, the Zeroth Law [8]. The laws are:

0. A robot may not injure humanity, or, through inaction, allow humanity to come to harm;
1. A robot may not injure a human being, or, through inaction, allow a human being to come to harm, unless this would violate the Zeroth Law of Robotics;
2. A robot must obey orders given it by human beings, except where such orders would conflict with the Zeroth or First Law;
3. A robot must protect its own existence as long as such protection does not conflict with the Zeroth, First, or Second Law.

These laws formed the basis for the theory of behavior, and governed human-robot interaction in his stories. At the heart of these laws is the robot's ability to recognize the human element and identify weakness or needs within humans. In promoting human-robot teams, the ability of the robot to possess at least a basic ability to recognize and evaluate human needs is essential.

Just as the human develops a theory of the robot's behavior, the robot must be able to understand and predict the human members of the team. This is not to say that machines must possess complex mental models or be able to discern our intentions. Rather, it is necessary to raise the level of interaction between the human and robot based upon readily available, non-intrusive workload cues emanating from the operator. The robot's theory of human behavior may be a rule set at a very simple level, or it may be a learned expectation developed through practiced evolutions with its human counterpart. The robot must possess some means to infer the need for intervention. Currently, accurate and non-intrusive collection of these signals is difficult at best, and those measures that have been used are unreliable at worst [9].

The answer to this dilemma is to reduce the human signals down to a prescribed set of channels, which are available as an integral part of the interaction of the human with the machine, and which the machine can use to configure its behavior and level of initiative.

Scientists in academia, industry, and government are currently researching methods that promote machine cognition of human needs. Key measures to assess human need and performance that researchers seek to identify include workload, attention, arousal, stress, fatigue, memory, and degraded performance. Current methods being evaluated in this area are listed in Table I. [10]

TABLE I
POTENTIAL HUMAN PERFORMANCE MEASURES

Neural Measures	Behavioral Measures	Physiological Measure
EEG	Eye tracking	Pupil dilation
ECOG	Reaction time	GSR
FMRI	Error rates	Heart rate
SPECT	Motion	Pulse
PET	Mouse pressure	Respiration
ERP	Key pressure	Temperature
MEG	Blink rate	Blood pressure
EROS	Control	Impedance
EMG	interactions	Saliva
EOG	Head position	Hormone levels
	Interface activity	Movement
	Vehicle data	Expired CO2
	Subjective reports	Oxygenation

While these measurements all present challenges, the utilization of human-robot teams confound the problem even more. While many of these measures may be assessed on someone sitting within a cockpit or behind a computer panel, they may not be measurable (at least at this stage in technology) for interaction with a mobile autonomous robot.

In the perceived human-robot teams, the robot is a situated and embodied entity that exists in the world potentially along side of the human element. Rodney Brooks, the Director of the MIT Artificial Intelligence Laboratory, provides the following explanation of these terms:

> A situated creature or robot is one that is embedded in the world, and which does not deal with abstract descriptions, but through its sensors with the here and now of the world, which directly influences the behavior of the creature.
>
> An embodied creature or robot is one that has a physical body and experiences the world, at least in part, directly through the influence of the world on that body. A more specialized type of embodiment occurs when the full extent of the creature is contained within that body. [11]

During the performance of a task, the human and robot may be side by side in the same proximal area or they may be in physically distal locations. Additionally, their proximity may change during the course of the task. Interaction between the robot and human may be through direct communications (verbal, gesture, touch, radio communications link) or indirect observation (physically struggling, erratic behavior, unexpected procedural deviation). Interaction may also be triggered by the observation of environmental / external factors (rising radiation levels, the approach of additional humans, etc.).

Once a prescriptive set of signals are identified, still further challenges remain such as how the machine should indicate its ability to do the task when the signals indicate that the human has attention elsewhere. It is critical that the machine disrupts the human only when absolutely necessary. In human only teams, one member can simply ask the other member if they require assistance. However, in mixed teams, it is critical to have cues to human performance that do not require special human attention. On the other hand, there will necessarily be situations where the robot must inform the human before taking action.

The robot's expectations must allow it to recognize human limitations and anticipate human needs without second-guessing the human's every move. When robots do intervene with their human counterparts, the human's TORB must be able to explain why the robot has stepped in and what this shift in control means for the task at hand.

To understand the need for a robot to possess a theory of human behavior, consider the case of a remote sensor deployment. In a remote setting, it is difficult though not impossible for the robot to recognize when the human needs assistance. By analyzing the commands sent by the human, the robot can infer human need. For example, if the operator frequently activates the robot's brake while the robot is traveling at high speeds, the robot may infer that the human is having difficulty driving the robot. Of course, this could be caused by several different factors including a lack of operator skill, a lack of operator confidence, high workload, or a very cluttered and/or dynamic environment. To disambiguate the possibilities, an effective theory of human behavior must draw from multiple cues. For instance, it is possible to monitor the frequency of commands issued by the operator. Another method we have implemented is to monitor the number of times that the robot is given a "dangerous" command. Dangerous commands can be recognized in terms of obstacle avoidance behaviors, tilt-sensors, bump sensors, inertial sensors, and resistance to motion measurements.

In addition to those outlined above, cues to operator workload could include: frequent pauses before initiating action; increases in errors or the time to perform a task; frequent re-orientation of the robot to alter the user's perspective on the display; repetitive behaviors; or simply informing the robot explicitly that the operator does not have the resources to devote to the task. Khoury and Kondraske [12] have also developed a reliable, non-intrusive means to assess human workload based on the interaction with the controller (e.g., joystick), which we hope to evaluate in the near future.

Additional performance measures / indicators (detailed in Table I) hold future promise, but may require the human operator to be augmented with instrumentation to provide the robot with the needed information. This is not inconceivable for near future implementation and is a needed area for research.

C. Robot Initiative

Once accurate measures of human performance are available to the robot, a new question arises: How should the robot initiate a new level of authority within the team? There are several ways the robot could do this.

1. The robot requests control of the task from the operator. This type of interaction could become frustrating if the robot is not able to assess accurately the ability of the human to perform the task and the relation of the performance measure to obtaining the goal. If the robot constantly pesters the operator about his or her performance, the operator would soon disable that ability.
2. The robot could simply state that it is taking control of the task, but allow the human operator veto power to this initiative. The system could become unusable or unwieldy, however, if the human has to frequently reverse actions taken by the robot or prevent the robot from taking an action.
3. The robot informs the operator it is taking control of the task, but denies veto power, if time is critical or the action to be taken has a small impact on performance.
4. The robot takes control of the task, and does not inform the human operator. In this case the operator may be incapacitated or delirious in which informing the human element serves no purpose or may have a detrimental effect.
5. The robot aborts of the task and returns to a base state. In this case, the robot possessing environmental sensors and an expectation of the human element may deem that the task is unachievable and termination is necessary to protect human life.

The benefits of allowing the team members to change roles within the team significantly increases team flexibility and reliability in task performance. However, if the interface and human-robot system are not designed in accordance with critical principles of human factors in mind, dynamic role changing may result in mode confusion, loss of operator situation awareness, loss of operator confidence in assuming supervisory control, and degraded and potentially catastrophic performance. Appropriate feedback is required when roles and levels of initiative change. Failure to inform the operator when the robot has overridden commands will lead to distrust of the system, unless the behavior is beneath the level of operator concern. This phenomenon has been seen in the airline industry with pilots and the automatic pilot mode of operation. [13]. One of the most significant elements in learning and developing system expectations is feedback.

The importance of feedback cannot be understated as the human and robots work within a mutual initiative system. Given that the human and robot may be in either proximal or distal environments, feedback may be provided via a computer based interface for interfacing with the robot or a visual or audible signal from the robot itself. Feedback from the robot should not only include the mode change, but also an indication of the reason for the change. For optimal performance of the team, the human must be able to develop expectations regarding when and why the robot will be motivated to initiate a new level of authority.

In fact, the level of initiative is of crucial importance. In order for the human's theory of system behavior to comprehend and exploit robot initiative, the robot's autonomy should be structured hierarchically such that at any given time, the user will know the bounds on what initiative the robot can take. Consequently, the INEEL has developed a control system with four clearly distinct levels of autonomy discussed in the following section.

IV. A CASE STUDY IN HUMAN-ROBOT INTERACTION

Robotic solutions are increasingly desired for conducting remote tasks in hazardous environments. For instance, remote characterization of high radiation environments is a pressing application area where robotic solutions can provide tremendous benefit. However, the DOE roadmap for Robotics and Intelligent Machines states that much more work is necessary in the area of human-robot interaction. In terms of time, cost, and safety, 'usability' is the most crucial component of robotic systems for remote characterization and handling of radioactive and hazardous materials.

In 2001, the INEEL utilized a robotic system coupled with a Gamma Locating Device (GLD) to characterize an area that had been closed to human entry for many years. This state of the art remote robotic system offered a means to remove the human from hazardous environments. However, the robot required high-fidelity video, reliable, continuous communication, and instrumentation of the environment *a priori*. As a mechanical 'subordinate,' this robot was dependent on continuous, low-level input from a human and was unable to cope with communication dropouts or continue the task when excesses in operator workload disrupted the input from the operator.

Fig. 1. INEEL operational robotic platform, an augmented ATRVJR.

In response, research efforts at the INEEL have developed a robotic system that can adjust its level of autonomy on the fly, leveraging its own, intrinsic intelligence to meet whatever level of control is available and appropriate. The resulting robotics system, pictured in Fig. 1., including hardware, software, and interface components, can slide between roles of 'subordinate,' 'equal' and 'leader.' For this system to meet its goals, we must be able to guarantee that the robot will protect itself and the environment. To do so we fuse a variety of range sensor information including inertial sensors, compass, wheel encoders, laser range finders, computer vision, thermal camera, infrared break beams, tilt sensors, bump sensors, sonar, and others. Also, a great deal of work has been focused on providing situation awareness to the user that can appropriately support the current level of interaction.

Our research to date has developed a control architecture that supports the following modes of remote intervention:

1. Teleoperation
2. Safe Mode
3. Shared Control
4. Full Autonomy

For each of these levels of autonomy, perceptual data is fused into a specialized interface that provides the user with abstracted auditory, graphical and textual representations of the environment and task that are appropriate for the current mode.

A. Teleoperation

We have taken the interaction substrate used in previous INEEL teleoperated robotic systems and revamped it through feedback with people who have deployed such systems. Within teleoperation mode, the user has full, continuous control of the robot at a low level. The robot takes no initiative except to stop once it recognizes that communications have failed.

B. Safe Mode

User directs movements of robot, but the robot takes initiative to protect itself. In doing so, this mode allows the user to issue motion commands with impunity, greatly accelerating the speed and confidence with which the user can accomplish remote tasks. The robot assesses its own status and surrounding environment to decide whether commands are safe. For example, the robot has excellent proprioception and will stop its motion just before a collision, placing minimal limits on the user to take the robot's immediate surroundings into account. The robot also continuously assesses the validity of its diverse sensor readings and communication capabilities. The robot will refuse to undertake a task if it does not have the ability (i.e., sufficient power or perceptual resources) to safely accomplish it.

C. Shared Control

The robot takes the initiative to choose its own path, responds autonomously to the environment, and works to accomplish local objectives. However, this initiative is primarily reactive rather than deliberative. In terms of navigation, the robot responds only to its local (~ 6-10 meter radius), sensed environment. Although the robot handles the low level navigation and obstacle avoidance, the user supplies intermittent input, often at the robot's request, to guide the robot in general directions. The problem of deciding how and when the robot should ask for help has been a major line of HRI enquiry and will be a major issue in our upcoming human subject experiments.

D. Full Autonomy

The robot performs global path planning to select its own routes, requiring no user input except high-level tasking such as "follow that target" or "search this area" (specified by drawing a circle around a given area on the map created by the robot). This map is built on the fly and uses frontier-based exploration and localization to perform searches over large areas including multiple rooms and corridors. The user interacts with the map to specify tasks and can guide the robot and infuse knowledge at an abstract level by selecting areas of interest and identifying sensed environmental features, which then become included within the map.

The latest development, and perhaps the most innovative aspect of our project to date, is that we have imparted a "theory of human behavior" within the robot's intrinsic intelligence, which allows the robot to assess human performance. Before we implemented this theory of human behavior, the robot was already able to use its knowledge of the environment and its own proprioception to take initiative and refuse to accept dangerous commands. However, the level of robot initiative was always controlled by the human. The "theory of human behavior" allows the robot to switch modes when the robot recognizes that the human is performing very poorly. This theory of human behavior is based primarily on the frequency of human input and the number and kind of dangerous commands issued by the user. For instance, if the human has repeatedly placed the robot or the environment in danger, or if the human has been unsuccessful in extricating a robot from a cluttered area, the robot will step in and take over from the operator. Although the human can ultimately override this capability, it provides a means for true peer-peer interaction.

V. CONCLUSIONS

This paper has presented some of the unique challenges associated with developing human-robot teams. Specifically it has explored the concept of mutual-initiative between human and robot team partners. Within this framework it is argued that it is essential that both the human and robot develop theories or expectations of how team members will react within a changing task environment. Paramount to this is the ability for the robot to recognize human needs and identify intervention points to take initiative.

The INEEL is currently exploring new ground in the area of human interaction with robots. Just as policemen develop intuitive relationships with their canine partner, so should be the relationships between human and robot team partners. The motivation for our work is the development of flexible human-robot teams to support the performance of tasks within human-hazardous environments.

Utilizing a robot equipped with robust sensors and intelligence, we are developing a human-robot control system and associated computer interfaces that promote mutual initiative between the human operator and the robot. Frequently operating in distal environments, the robot is often able to make better judgments about its environment (i.e., local navigation) than humans. Consequently, we have created modes of control where the robot monitors human command input and infers the need to supplement or override human action. The robot has the power to refuse to undertake commands from the user that are deemed by the robot to pose a threat to itself or its environment. This engenders a host of new questions, especially in regard to how an autonomous and mobile robot can infer intervention points. Our current system utilizes human command sequences as a measure of human performance. Within our implementation, human error loses much of its sting because the robot is able to countermand dangerous commands. At the same time, we have provided means for the human to override robot initiative and to configure the robotic initiative for specific tasks and environments. In this way, the human and robot become true team partners who can support and compensate for one another to adapt to new challenges.

Although this paper explores a new paradigm for human – robot teaming, it must be acknowledged that our current mutual-initiative robotic system presents only first steps towards the possibility of robots assuming peer or even leadership roles. Future research will focus on the development and evaluation of new approaches to provide greater precision in determining the human state and performance level.

VI. REFERENCES

[1] Sheridan, Thomas B., "Supervisory Control of Remote Manipulators, Vehicles and Dynamic Processes: Experiments in Command and Display Aiding." In *Advances in Man-Machine Systems Research*, vol. 1, W. Rouse, Ed. Greenwich, Connecticut: JAI Press Inc., 1984, p. 50.

[2] Sheridan, Thomas. B.. "Supervisory control," in *Handbook of human factors* (2nd ed.) G. Salvendy, Ed. New York: Wiley, 1997. pp. 1295-1327.

[3] Sheridan, T. B. & Verplank, W. L. "Human and computer control of undersea teleoperators," Man-Machine Systems Laboratory, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge. Mass., Tech. Rep., 1978.

[4] Yen, J., Yin, J., Ioerger, T., Miller, M., Xu, D., & Volz, R. "CAST: Collaborative Agents for Simulating Teamwork.," In *Proc. of the 17th Int. Joint Conf. on Artificial Intelligence*, Seattle, WA, 2001.

[5] Royal Canadian Mounted Police: Police Dog Services, "K-9 Recruit Successfully Completes Training," 2002. [Online]. Available: http://www.rcmp-grc.gc.ca/html/dogs.htm

[6] Brooks, R., Brazeal, C., Marjanovic, M., Scassellati, B., & Williamson, M. "The Cog Project: Building a Humanoid Robot." In *Computation for Metaphors, Analogy, and Agents, Springer Lecture Notes in Computer Science, 1562*, C. Nehaniv, Ed., Springer-Verlag, Berlin, 1998.

[7] Asimov, I. *I, Robot (a collection of short stories originally published between 1940 and 1950)*, London: Grafton Books, 1968.

[8] Asimov, I., *Robots and Empire*, London: Grafton Books, 1985.

[9] Wickens, C. D., *Engineering Psychology and Human Performance*, (2nd ed.), NY, NY: Harpers Collins Publishers, Inc. 1992.

[10] Lyons, D., "Cognitive Workload Assessment." In *Proc. 2002 DARPA Augmented Cognition Conference*, pp. 64.

[11] Brooks, R. A., *Flesh and Machines*, New York: Pantheon Books, 2002, pp. 51-52.

[12] Khoury, G. J., & Kondraske, G. V. "Measurement and continuous monitoring of human workload associated with manual control devices," NASA JSC;. Arlington, TX: University of Texas at Arlington, Human Performance Institute, Tech. Rep. 91-011R, 1991.

[13] Abbott, K. A., S. M. Slotte, D. K. Stimson, "Federal Aviation Administration Human Factors Team Report on: The Interfaces Between Flightcrews and Modern Flight Deck Systems," Federal Aviation Administration, Washington, DC, Tech. Rep. June, 1996. [Online]. Available: http://www.faa.gov/avr/afs/interfac.pdf, pp. D1-D3.

V. BIOGRAPHIES

David J. Bruemmer received his B.A. in Computer Science and a B.A. in Religious Studies from Swarthmore College. He has worked as a technical advisor for robotics programs at the Defense Advanced Research Projects Agency (DARPA) and as Principal Investigator on a DARPA project to develop an autonomous landmine detection robot. He is currently the Principal Investigator for a research project investigating advanced human-robot interaction and adjustable robot autonomy. His research interests include autonomous agents, distributed robotics, adaptive behavior, and human-robot interaction.

Dr. Julie L. Marble received her doctorate in Cognitive Psychology/Human Factors from Purdue University. Before joining the INEEL, Dr. Marble worked in the Usability Test Center of QWEST; conducting usability analysis and testing, interface design, and performance measurement of software systems. Her research on human performance includes the effects of display designs for teleoperation of robotic systems under normal and failure conditions, interactions of displays with control systems, and display design on human performance. Her other research has included human error in aviation maintenance, spatial decision-making, and serial implicit spatial learning.

Donald D. Dudenhoeffer received his Masters of Science degree in Operations Research from the Naval Postgraduate School in 1994. He has been a Principal Investigator for numerous project in the area of robotic, AI, and human-robot interaction including: Behavioral And Control Modeling for Large Scale Micro-Robot Systems (1999-2002), Command and Control Architectures For Autonomous Micro-Robotic Forces (2000-2002), and Integration of Artificial Intelligence Systems for Monitoring And Control of Complex Process Systems (2000). His research interests include simulation, robotics, human performance modeling, military operations and command and control. Prior to working at the INEEL, he served 12 years on active duty as a submarine officer and currently holds the rank of LCDR in the U.S. Navy Reserves.